

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo 426

August 1977

**THE COMPUTATION OF
IMMEDIATE TEXTURE DISCRIMINATION**

Bruce R. Schatz

Short abstract

The computation of immediate texture discrimination involves finding boundaries between regions of differing texture. Various textures are examined to investigate the factors determining discrimination in the limited domain of line-and-point images. Two operators embodying necessary properties are proposed: length and orientation of actual lines and of local virtual lines between terminators. It is conjectured that these are sufficient as well.

Relations between this theory and those of Julesz and of Marr are discussed. Supporting psychological evidence is introduced and an implementation strategy outlined.

Revised version of a thesis submitted in May 1977 to the Department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology in partial fulfillment for the degree of Master of Science.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research Contract N00014-75-C-0643.

ABSTRACT

The computation of immediate texture discrimination involves finding boundaries between regions of differing texture. The properties essential to determining this are investigated here, and two operators given which together appear necessary and possibly sufficient to accomplish this in a limited domain.

A psychological definition for discrimination is given. In the simple case of texture elements composed of lines and points regularly repeated over a large matrix, a number of possible properties are examined for effects on discrimination. Two operators, embodying necessary properties, are proposed:

- (1) length and orientation of lines
- (2) length and orientation of local virtual lines .

Virtual lines are imaginary lines which behave as though physically present. They are drawn between terminators (line ends) here. Evidence for the necessity of lines and each type of virtual line is supported by example textures.

Marr proposed that first order distinctions are sufficient; necessity of the second operator provides an improved lower bound on what is necessary to compute texture discrimination. Julesz conjectured that two textures are not discriminable if their second order statistics are identical; the computation here, shown to be strictly less powerful and using a proper subset of the dipoles, is proposed to provide an improved upper bound as well.

Psychological evidence for the reality of the assumptions behind the computation is introduced and extensions to less restricted domains discussed.

An implementation strategy is described. This basically consists of moving a window over the image and comparing the orientation and length changes in adjacent windows. When a sufficiently large change is found, a texture boundary is asserted. Details of the window comparison and parameter values are given.

Thesis supervisor: David Marr,
Research Associate, Dept of Electrical Engineering and Computer Science

Acknowledgements

My thesis supervisor, David Marr, provided much of the philosophy behind this work as well as a myriad of helpful suggestions and careful criticism.

Shimon Ullman gave more than could have possibly been asked for, from perspicacious suggestions, both major and minor, to philosophical comment to moral support. I remain deeply in his debt.

In the earlier days of this research, Tomas Lozano-Perez also provided that type of consistently wise advice that one is eternally grateful for.

Berthold Horn, David Marr, and Shimon Ullman read manuscript drafts; their constructive criticism substantially improved both the form and content.

Whitman Richards added greatly to my understanding of psychological results.
Bob Woodham suffered through several discussions of computational issues.
Keith Nishihara and Kent Stevens made a number of useful suggestions.

The uniquely sage and flexible computing facilities at the M.I.T. Artificial Intelligence Laboratory have provided much assistance and many pleasurable hours.

Finally I thank the National Science Foundation for a graduate fellowship which provided financial support throughout most of this research.

The author's current address is:

Computer Science Dept, Carnegie-Mellon University, Pittsburgh, PA 15213.

Table of Contents

Chapter 1: Immediate Texture Discrimination

1.1 Texture and its computation	5
1.2 Background	10
1.3 Experimental conditions	18

Chapter 2: Essential Determining Properties

2.1 Introduction	21
2.2 Local shape operators	23
2.3 Virtual lines and terminators	32
2.4 Relation to Julesz's theory	46
2.5 Psychological support	52
2.6 Extensions	56
2.7 Other possible operators	62

Chapter 3: Computation of Texture Boundaries

3.1 Overview	71
3.2 Windows	72
3.3 Comparison of windows	75
3.4 Parameter values	86
3.5 Conclusion and future prospects	91

Notes	93
-------	----

References	94
------------	----

Chapter 1 Immediate Texture Discrimination

1.1 Texture and its computation

The computation of immediate texture discrimination involves finding boundaries between regions of differing texture. But what features of the image must be examined and what operators should utilize these? This thesis considers these questions in a restricted domain with texture elements composed of lines and points. A psychological definition of texture is set up and various properties investigated. Evidence is given that two proposed operators are necessary to characterize discrimination, and it is conjectured that they are sufficient as well. Thus an improved lower bound (from Marr) is provided for the texture computation and an improved upper bound (from Julesz) conjectured. An implementation strategy is also discussed. [Note 1]

The first step is to define the problem to be considered. Texture, according to the *Oxford English Dictionary*, is the constitution, structure, or substance of anything with regard to its constituents or formative elements. It can also be regarded as the structural property which makes surfaces appear as surfaces as opposed to insubstantial areas (Gibson[1950, p53]). Some objects do not possess texture, plate glass or steel balls for example; these will not be considered. In the visual world, texture is the result of perception of physical irregularities of surfaces such as bumps, weaves, dips, or graininess,

or perception of changes in reflectivity. In the intensity image, this corresponds to considering the constituents comprising a region as opposed to the region as a whole. A texture region is thus a collection of proximate, similar elements. The formation of texture crucially depends on small, fairly similar elements repeated fairly regularly over a large area. (Hawkins[1970],Pickett[1968]). The question examined here is how to predict texture discrimination, how to separate regions on the basis of immediate texture differences. Problems of identifying or classifying textures are not considered. "Texture discrimination" will subsequently refer to *immediate* discrimination of differing regions.

As natural textures can be quite complicated (c.f. Brodatz[1966]), artificial examples will be used to examine discrimination. These will consist of a simple element, composed of lines and points, regularly repeated (with small perturbations) over a large matrix, thus generating a texture. There will be two regions, one inside the other, with one generating element per region. The goal is to find operators which determine whether the inner region is perceptually distinguishable and if so, compute its boundaries. (See fig 1.1.1.) Immediate perceptual discrimination is determined experimentally by the class of texture pairs for which the differing region can be identified in 200 milliseconds (disallowing scrutiny by directed eye movements -- see section 1.3). The line and point restriction yields a natural but non-trivial subset of possible images.

The effects on texture discrimination considered here will be those directly related to properties of the generating elements. Overall region properties, including those derived from averaging dissimilar elements over the entire region, will *not* be considered. Examples of these measures, which will be held constant, are overall brightness (region A is darker




Figure 1.1.1.1

than B), color (region A is red, B blue), and motion (region A is moving differently than B). In other words, two regions being compared for possible texture differences will be assumed to have the same overall brightness (average intensity), the same color (black/white), and the same motion (none). In addition, all the processing is assumed to take place in a flat two-dimensional world before depth information is used. Spatial frequency differences will be ignored by assuming constant density of points, and the spacing between elements will be assumed large enough to avoid interaction effects between them. Accordingly, discrimination can be determined by comparing the generating element of adjacent textures. Since this is the case, a texture and the element which generates it will be referred to interchangeably, often mentioning only the element.

A methodological point should be mentioned. Most of the psychological literature on texture is concerned with experimental results with little or no mention of computational issues, while most of the engineering literature discusses operators useful for specific problem domains. (Compare, for example, Pickett[1970] to Hawkins[1970].) In contrast, the focus here will be directly on the computational *problem* of texture discrimination and an attempt made to discover the psychologically essential properties affecting it along with a feasible computation for determining the boundaries of immediate texture regions. Thus (informal) psychophysical experimentation will be used as the criteria for inclusion of properties while computer simulation will be used to test the computational effectiveness. Note that even if one is primarily interested in the construction of machine perception, humans provide a good instance of effective visual processors and thus it may be reasonable to imitate them. Similarly, machine vision can provide much information about

the basic structure of biological visual processors since one is required to consider potential mechanisms in detail.

The structure of the thesis is as follows:

The remainder of Chapter 1 describes background material, primarily the work of Julesz and Marr, as well as the experimental conditions used in viewing the textures.

Chapter 2 introduces the operators, which are derived from properties that appear necessary and possibly sufficient to predict texture discrimination of the type considered. These are the length and orientation of actual lines and of local virtual lines between terminators (imaginary lines filled-in between special points.) Evidence for the necessity of actual lines and for each type of virtual line is provided. The necessity of using virtual lines produces an improved lower bound on what is needed for the computation. Julesz conjectured that two textures are not discriminable if they possess the same second order statistics. (This occurs when the length and orientation distribution of all possible line connections (dipoles) between points is the same.) The operators here use a proper subset of the dipoles to provide a feasible computation, strictly less powerful than second order statistics, which is tentatively proposed to be sufficient; thus producing an improved upper bound from Julesz's. Extensions to less restricted domains are considered, e.g. ideas on handling proximity effects caused by spacing. Supporting evidence from the psychological literature is surveyed. Other possible operators are examined.

Chapter 3 discusses details of the computation. This basically consists of moving a window over the image and comparing the orientation and length changes in adjacent windows.

When a sufficiently large change is found, a texture boundary is asserted. Descriptions of the window comparison and the parameter values are given.

1.2 Background

There are two major predecessors to the approach to texture discrimination outlined in this thesis. Julesz [1962,1965,1973,1975] developed the paradigm of examining texture discrimination by considering a small matrix generated by one simple element inside a large matrix generated by another and asking whether the inner matrix could be distinguished in short periods of time. His conjecture of a second order statistical limitation on discrimination was very insightful; it provides a small upper bound on what needs to be used in a feasible texture boundary computation. Marr[1976] argued for a process-oriented explanation of texture vision; he gives a set of primitives, representing the first stage in the vision process, and conjectures that first order discriminations on these are sufficient. The examples given here demonstrate that the lower bound is in fact greater than first order and that the upper bound is probably lower than second order. Thus the power necessary lies somewhere in-between Marr's and Julesz's proposals. These predecessors will now be explained in further detail followed by a short survey of other work in texture.

Many researchers have argued for processing texture directly from the intensity array

(e.g. Rosenfeld[1971], Haralick[1973]). The point of view taken here is different, i.e. only a processed version of the image is used for input. Marr[1976] makes a strong argument in favor of this and his views are concurred with here. One reason is that humans appear to perform perceptual processing on symbolic descriptions of an image rather than on the direct intensities. For example, in a Cornsweet edge (Ratliff[1972]), two adjacent regions of identical intensity appear different because an edge is perceived between them. The primitives propounded by Marr, the *primal sketch*, consist of edges, lines, and blobs. These are described by their orientation, length or size, position, termination points, and local contrast. The concern here will be with elements composed of lines and points; possible extensions are discussed in section 2.6.

Marr proposed that first order discriminations on these primitives are sufficient to account for texture discrimination. Strictly speaking, the existence of discriminable textures generated by elements differing solely in virtual lines shows this to be false. In terms of primal sketch properties, the orientation and length of lines are in fact used but the position is an absolute not a relative value. Thus with first order distinctions (one-dimensional histograms), arrangement differences, where the same actual lines appear in different relation to each other, cannot be detected although they can cause discrimination (c.f. fig 2.2.5). However, by connecting terminators (somewhat differently defined), virtual lines can be generated and first order operations on these currently appear to be sufficient. Thus the verdict on Marr's proposal depends on whether the formation of virtual lines are included in his processing or not. Connecting points is not a first order operation (it is of second order) but his claim was that grouping processes are included as well. If the proper virtual

lines are formed before the first order processing, then his proposal is correct. Perhaps the best point of view is to not consider order statistics but instead discuss what properties are necessary and what operations should be performed on these. Marr's proposal then becomes a highly suggestive sketch of a computational theory which is lacking many essential details.

The most comprehensive proposal regarding texture discrimination thus far is that of Julesz [1973,1975]. He conjectured that two textures are not discriminable if their second order statistics are the same. This occurs when for each possible length and orientation, there are the same number of dipoles of that length and orientation in each texture. A dipole is a possible line segment connecting two black points of the image. Dipoles can thus be connected between any two points within a physically present line segment (including subsegments) and between any two points each of which is either on a line or is an isolated point (joining two points not physically connected). So from one point on a line there are an infinite number of dipoles between it and (the points on) another line; these are of varying lengths and orientations. [Note 2] Consider fig 1.3.1, assuming the (maximal) line segments are of length 1. The two textures have identical second order statistics; each generating element has an infinite number of dipoles of length 1 at 0 degrees (including one physically present), two dipoles of length 1 at 90 degrees, an infinite number of dipoles at 0 and at 90 degrees of varying lengths less than 1, and an infinite number of dipoles connecting the lines at varying orientations and lengths. The prediction is accordingly that the textures are indiscriminable, as they in fact are. Julesz, et.al. give three methods for constructing, from a given generating element, another element which can generate a texture




Figure 1.3.1

with the same second order statistics. One such method is the one illustrated here, namely rotating it 180 degrees. (Gilbert and Shepp[1974] prove that the methods have the desired property.)

It may be helpful to view order statistics in the following manner. First order statistics measure the probability f_0 that a (black) point appears at any given coordinate. For two textures to have the same second order statistics, they must have the same f_0 and in addition, the same conditional probability that a line of a particular length and orientation occurs between a pair of points. Thus the probability $f(r)$ that the randomly placed vector \underline{r} touches a point on each end must be the same in both textures for all \underline{r} . A simple example where this is not true is two random dot patterns with identical f_0 but one of which has the property that the minimum nearest neighbor distance is 10. The dipoles are those vectors \underline{r} which do touch (black) points on either end.

The converse of the above conjecture is often but not always not true. For example, fig 1.3.2 is a case where the statistics differ (considering dipoles of length 1: the outer generating element has two vertical and many horizontal while the inner has two horizontal and many vertical) and the textures are in fact discriminable. However, in fig 1.3.3, the statistics differ yet the textures do not. Thus the claim is that second order statistics being identical suffices to predict indiscriminability of textures (but is not necessary). In this thesis, Julesz's conjecture will be assumed to be correct and thus provide an upper bound limitation. The proposed operators use a subset of the dipoles to provide a less powerful computation which may still be sufficient.




Figure 1.3.2




Figure 1.3.3

For a survey of the engineering literature on texture see Hawkins[1970]. For a survey of the psychological literature see Pickett[1968,1970]. Lipkin[1969] contains a bit of both. Riseman & Arbib[1977] give more recent material. A common division in the computing literature is between structural (placement rules on a unit generating pattern) and statistical (numeric averaging descriptions of local properties of regions) approaches (e.g. see Zucker[1976(a)]). The approach here is a hybrid as it performs pseudo-statistical analysis on unit generating patterns.

A common approach which ought to be sketched here consists of taking a two-dimensional Fourier transform of the image and looking for differences in gross measures on the power spectrum, which can often discriminate between texture regions (see e.g. Bajcsy[1973], Lieberman[1974]). Common measures include directionality and element size and spacing; these can be computed as follows. Let $P(r,\theta)$ be the power spectrum of the image in polar coordinates. Let $P(r)$ be the sum over θ for given r and $P(\theta)$ the sum over r for given θ . The directionality is determined from the sharp peak (if any) in $P(\theta)$. Since the power spectrum is invariant with respect to translation and a texture consists of regularly repeated elements, the period can be determined from the maximum $P(r)$. Then the element size is the width in r of that peak, and the spacing is the period minus the element size. These measures are much less sensitive than those proposed below. This has the result that they would often fail to distinguish between discriminable textures, including several here, with the same spatial frequency and the same element size. However, they can deal with many types of textures not investigated here.

Many of the texture elements were adapted from Julesz or Beck.

1.3 Experimental conditions

This section contains a description of the experimental conditions under which discrimination was tested.

As mentioned above, the textures consist of a small matrix of one element inside a large matrix of another and the question asked is whether the inner region can be distinguished immediately. But what does "immediate" mean in this low-level processing task? Julesz spends much of his time intuitively discussing the issue. Presumably what is desired is to consider only pure perception, i.e. allow no cognitive processing. In early stages of this research, it became clear to the author that informal presentation of textures was not producing consistent results (e.g. there occasionally was strong disagreement about whether discrimination occurred or not). So a more precise delineation of the class of textures to be examined became necessary.

A reasonable interpretation of pure perception is to disallow scrutiny. This means not allowing the viewer to successively focus in several places while carefully looking for differences. If one allows an indefinite length of time, then clearly any two non-identical textures can be discriminated. So assuming that scrutiny implies directed refocusing of attention, the restriction of allowing only a few eye movements (saccades) was made.

However, there still remains the problem of whether the time limit chosen is arbitrary. For instance, is the class of textures discriminable allowing one saccade significantly different from that when two are allowed? If so, the concept of pure perception is rather suspect and one might as well cease attempts at explaining texture discrimination (since allowing unlimited time requires operators which can distinguish between arbitrarily similar but non-identical textures.)

To check this, a number of texture pairs ranging from obviously distinct, fig 1.1.1, to apparently identical, fig 1.3.1, were viewed for 200 milliseconds and for 2000 milliseconds. (200 ms is not sufficient time to allow a saccade driven by the stimulus while 2000 allows several such). For the most part, the two textures were either discriminable in 200 ms or still not discriminable after 2000. Thus it appears reasonable to suppose that there is a fairly distinct class of textures discriminable using only low-level pure perception and it is this class this paper seeks to compute.

At this point, it is worth noting the precise experimental procedure used in testing the texture discriminations.

A large matrix of element A was presented with a smaller matrix of element B contained within it. (Typical sizes were 12 by 12 and 3 by 6; the elements subtend an angle of about 0.3 degrees at 3 feet away from the viewer.) In most experiments, the intensity of points was constant (either black or white). The elements were regularly repeated over the matrix. In most of the textures, all elements had a random jiggle to prevent extraneous columnar effects from influencing discriminability. A single element was used in each texture (no

variation). The spacing in both x and y was varied in different experiments, but was constant throughout any one texture. There was a fixation point in the center.

Basically the procedure was to flash a texture on a CRT screen (controlled by a minicomputer; DEC GT44 attached to PDP 11/40) for 200 ms and then mask off further visual input by displaying a full screen of random dots. Subjects were told they would be shown two textures and were asked if they could discriminate between them. If the answer was yes, they were asked about the differing subregion:

- (a) its location (lower left or upper right), and
- (b) its shape (vertical rectangle or horizontal rectangle).

The questions (forced choice) were designed to insure that the differing region, and not an accidental view of two adjacent but differing elements, was the cause of the perceived discrimination.

Two textures were classified as discriminable if both questions were correctly answered by a significant proportion of the subjects. Thus immediate region discrimination was being tested. This procedure produced reasonably consistent results which agree well with the intuitive answers. However, thus far no full-scale rigorous experiments have been run although many textures, including all of the doubtful discrimination cases mentioned in this thesis, have been tested on the author.

Chapter 2 Essential Determining Properties

2.1 Introduction

This chapter will discuss properties essential for determining immediate texture discrimination and their embodiment in computational operators. As element-only comparisons are being considered here, the quest will be to find measures for comparing the generating element from adjacent textures, i.e. approximating its local shape. Note that two elements having similar shape in the texture case (when they each are repeated over a large area and the areas are compared) does not imply that they would be judged similar when directly compared to each other. In fact, with most of the indiscriminable texture pairs examined, the generating elements appear different when compared to each other. Remember also that only crude distinctions can be made in 200 milliseconds.

In the line-and-point world here, what operators should be chosen? Lines have 3 basic properties: length, orientation, and position. The first two are used directly while position is accounted for by postulating the existence of *virtual* lines (imaginary lines between points which behave as though physically present). Thus the two operators proposed are:

The length and orientation of

(1) actual lines and

(2) local virtual lines between *terminators*.

The terminators are isolated points, endpoints of lines, and corners. The comparison will be

made on adjacent elements by

$$(\sum \text{Differences}) / (\sum \text{Total}) > \text{Threshold}$$

where the differences and total refer to the length of lines at each orientation. The orientations are considered in groups of equivalence classes (buckets) and crude histograms of the lengths in each bucket are compared. (Note that this ratio means that the difference must be sufficiently great for discrimination to take place.)

Examples will be given of the utilization of actual and virtual lines; the corresponding operators are proposed to form a lower bound on what is needed for computing texture discrimination. The relation to second order statistics will be considered; it will be shown that the operators are strictly less powerful and conjectured that they are sufficient nonetheless. If true, this places separation of immediate texture regions after an initial processing of the image but before even a simple grouping into texture elements (c.f. section 2.7). Discussion of extension to less restricted domains and some psychological evidence is also given.

A fuller description of the operators and the rationale behind them follows. As above, the properties and the operators which embody them will often be referred to interchangeably.

2.2 Local shape

The situation is thus: two elements generating textures must be compared -- these are composed of lines and points. What features of the local shape are crucial in immediate texture discrimination?

Lines have three basic properties: length, orientation, and position. (Points can be considered as the limiting case where the length is zero so that only position matters.) Each of these will be considered in turn, and two operators proposed to account for them.

Orientation differences can cause strong discrimination. See, for instance, fig 1.1.1. In fact, two lines are distinguishable (in the texture case) when their orientations differ by more than about 10 degrees (see section 3.4). Thus the orientations will be considered in equivalence classes of ± 10 degrees. There is also psychological and neurophysiological data supporting the presence and importance of orientation sensitivity. This is discussed in section 2.5.

Length must be considered as well since gross differences in length can cause discrimination even when the orientations are the same. See fig 2.2.1 for instance. (The lines in the inner texture are twice as long as those in the outer.) Within a particular orientation class, the length comparison is done by a histogram (as opposed to the sum total). For example, in fig 2.2.2 with respect to lines of length 1, the outer element has 2 at 0 degrees and 1 at 90 while the inner has 1 at 0 and 2 at 90. In fig 2.2.3, both (a) and (b) have the same horizontal lines and the same vertical line sum. However, in (a), the vertical lines are nearly identical in each of the generating elements and there is no discrimination





Figure 2.2.1

[illegible]

Figure 2.2.2

INSIDE

**BAR-LENGTH: 23.
ORIENTATION:**



OUTSIDE

**BAR-LENGTH: 51.
ORIENTATION:**




Figure 2.2.4

whereas in (b), the vertical histograms now differ (1 long for the outer element versus 2 short for the inner) and discrimination occurs. In fig 2.2.4, the inside and outside textures have roughly the same line length density, yet discrimination occurs because of differing length distribution. (The inside has twice as many lines which are half as long.)

Since length at each orientation affects discrimination, the first operator which appears to be necessary is

(1) length and orientation of lines.

Note, as mentioned above, the length is compared via a crude histogram for each of the several small equivalence classes of orientation. Thus a two-dimensional histogram is being used.

Line here means maximal physically present line (subsegments ignored). No predictive power is lost by this restriction since two maximal lines differ if and only if some of their subsegments do. Lines are delimited by terminators (see below) so that a plus is considered to contain 4 lines.

Effects dealing with position must still be considered. Gross positional differences cannot occur since an element-only comparison is being dealt with. However, a number of discriminable textures, such as that in fig 2.2.5, are identical under operator (1), i.e. they have the same orientation and length of lines. The discrimination appears to be due to the fact that the lines are arranged differently. As suggested by such textures as fig 2.2.6, one solution is to introduce *virtual* lines, imaginary lines connecting points which behave as though physically present, although somewhat more weakly. These virtual lines are




Figure 2.2.5




Figure 2.2.6

intended to capture the local positional geometry of the image, in particular the relationships of lines and points to each other. It should be noted that virtual lines are not "necessary" in the sense of the actual lines of the first operators, that is, they are at present only a convenient fiction. However, they successfully explain a number of puzzling textures, they fit naturally into Marr's and Julesz's schemes, and they have several other desirable side-effects (as will be brought out). Thus the second proposed operator will essentially be operator (1) on virtual lines. However, there are a few points to be considered before stating it.

There are a potentially infinite number of virtual lines, and as the goal here is a feasible computation, these must be restricted to a small finite number. The restriction will be made by only connecting special points of some kind. Thus the question is which to choose.

The first problem is that, theoretically, any pair of points in the image could be considered as connected by a virtual line. To limit this, only *local* virtual lines will be drawn, i.e. for any special point, only special points in some small neighborhood of it will be considered for connection. (See section 3.4 for details.) This is actually more general than the class of textures considering only comparisons between elements. For if the spacing between elements happens to be small, then virtual lines between elements are drawn. This often successfully predicts a class of discriminations called "boundary effects" (see section 2.6).

Even within an element, there are usually a computationally impossible number of virtual lines. For example, between any two nearby lines there are an infinite number of possible connections. As the world here consists of lines and points, one obvious class of

places to connect virtual lines to are the *terminators* ("endpoints") of lines. In other words, only virtual lines connecting the ends of lines are drawn; it is claimed that these are the psychologically important ones. (Isolated points are also considered as terminators.) The next section explores the various types of terminators and demonstrates their "necessity" by examples which seem to require them. Thus the second necessary operator is

(2) length and orientation of local virtual lines between terminators.

The above operators are at least necessary to predict texture discrimination in the specified class if one believes the examples given. It is conjectured that they are sufficient as well; this is discussed in section 2.4. The operators are used to essentially compare adjacent texture elements. If these differ, i.e. $(\Sigma \text{difference}) / (\Sigma \text{total}) > \text{threshold}$, a texture boundary is asserted. This ratio, rather than a pure difference, is required to insure that the elements are not merely different but sufficiently different (see sections 2.4 and 3.3).

The next section considers cases in which positioning of the lines appears to determine discrimination, i.e. in which virtual lines between terminators play a major role.

2.3 Virtual lines and terminators

In a sense, virtual lines, imaginary lines which behave as though present, are not strange at all. They perform the basic task of representing the local geometry in the image, i.e. the relations of lines and points to each other. The necessity of this function has been

recognized since the days of the Gestaltists who discussed the strong effects of closure and continuation on shape perception. This section will give evidence of the uses of virtual lines and where they occur.

Virtual lines have also been introduced by other researchers as well. Attneave[1974,p116ff] mentions their use in apparent motion and refers to several others who found them convenient. Ullman[1977] found cases in motion of two points moving together as though linked. Stevens[1977] bases his computation predicting parallel structure in random-dot inference patterns on filling-in local virtual lines between points.

It was noted earlier that virtual lines will be created here only by connecting terminators locally. These are points of special psychological interest. Thus the question arises as to what these are. The following informal definition will be used.

Definition: A terminator is a point in a series of connected line segments where the slope changes significantly after having remained constant for a sufficiently long period.

So if one had a long straight line, there would be a terminator at the end. If a small curved piece was added to the end, there would still be a termination in the same place. Note this means that only one of the lines connected to a point need consider it a terminator for the point to be a terminator. This is further discussed with regard to smooth curves below. There are three types of terminators: isolated points, endpoints of lines, and "corners". Examples of these will be given in instances where they are logically necessary, i.e. where the actual lines in the elements under consideration are identical yet the generated textures differ. (This is a class of textures not recognized, for example, by Marr[1976].)




Figure 2.3.1





Figure 2.3.2

Isolated points are not strictly terminators under the definition, but there are cases in which virtual lines connecting them to lines or other points are needed in the scheme here. Examples of discriminations predicted by virtual lines between points are fig 2.2.6 and fig 2.3.1 . Note that since operator (1) on actual lines has no effect on isolated points, virtual lines connecting them are needed to measure their contribution to local shape.

Endpoints of lines are the prototypical termination points, and virtual lines which connect these provide a basic positioning capability. Generating elements differentiated solely by this type of terminator produce *arrangement* differences, i.e. the same actual lines occur in both elements but are positioned differently and the textures differ as a result. Examples are fig 2.2.5 where the outer texture element has a virtual line at 135 degrees while the inner has one at 45, and fig 2.3.2 where the outer texture element has a short virtual line at 0 degrees and the inner has a long one at 90. Virtual line differences are in general weaker than actual line differences. There is psychological evidence supporting this, which is discussed in section 2.5.

Corners are much more ambiguous than endpoints as one encounters murky areas of smooth curves and non-discriminable elements with different full dipoles. While the author has not found an example strictly requiring virtual lines between corners, they often provide useful explanatory power. Consider the following series of textures illustrated in figs 2.3.4 - 2.3.7 in which virtual lines are drawn between sharp corners (here right angled). There are three virtual lines within each element: (1) connecting the two endpoints, and (2) and (3) connecting the corners. As shown in fig 2.3.3, the endpoints (1) change from diagonals to vertical lines while the corners (2) change from vertical lines to diagonals (corners (3) do not



	$h=0$	$h=s/6$	$h=s/2$	$h=s$
endpt 1				
corner 2	none			
corner 3				
Discrimination Case	yes I	yes II	no III	no IV

Virtual line comparison ($s=30$)									
II					III				
	deg	length	deg	length		deg	length	deg	length
1	130	39	50	39		116	37	64	37
2	81	30	99	30		64	32	116	32
3	135	42	45	42		135	42	45	42
Σdif	$2*39 + 2*42$				Σdif	$2*42$			
$\Sigma total$	$2*39 + 2*30 + 2*42$				$\Sigma total$	$2*37 + 2*32 + 2*42$			
	.73					.38			

Figure 2.3.3 Virtual lines between corners




Figure 2.3.4




Figure 2.3.5




Figure 2.3.6








Figure 2.3.7

change). With small upper bars, the only match between the two elements is $A_2=B_2$ (case II) and discrimination occurs. (The difference ratio is .73 while the threshold is approximately .5 .) As the upper bar becomes longer, the virtual lines change. When it reaches the halfway point (is half the length of the lower; case III), A_1 now equals B_2 and A_2 now equals B_1 (in the sense that the orientations are now within 10 degrees of each other and the lengths are closely comparable). Thus, although the virtual lines (and by implication, the dipoles) are different, they are not sufficiently different to cause discrimination. (The difference ratio is only .38 .) This is an explanation as to why Julesz's conjecture of second order statistics is not necessary since it is a case in which the dipoles differ but the textures do not (c.f. fig 1.3.3) . It shows that the dipoles must differ by a sufficiently large (relative) amount to insure discrimination. However, unless one has only a small finite class of dipoles to consider and a specified mechanism for comparing them, determining this amount can prove very elusive. This is one reason as to why a computational theory, such as proposed in this thesis, may be preferable to a statistical one, such as Julesz's.

The above is an illustration of the use of sharp corners as termination points. But what if the corners are instead rounded so that the elements are now continuous smooth curves? (See fig 2.3.8 .) A problem with what should be considered termination points now arises. While in the sharp corner case one might be willing to consider the elements as composed of three lines with terminations occurring at the joins, in the rounded corner case the inclination is to say that there is only a single (curved) line (c.f. Koffka[1935,p151]). However, the shape is nearly identical in both cases and thus the explanations of discrimination should be nearly identical as well. Assuming lines to be only straight and



sharp corners



smooth corners

Figure 2.3.8 Termination points for corners.

applying the definition given in this section, termination points are placed in correspondingly same places in both instances. In either case, the large horizontal lines generate terminators at their ends, i.e. at the point where the long constant slope has now changed. However, in the sharp case the vertical lines also produce a termination assertion (which coincides with the place asserted by the horizontal). In the rounded case, the smooth arc which forms the vertical piece consists of small segments with smoothly changing slope and thus no terminators are asserted by these. Still, in this case terminators appear at the correct points since they are asserted by the horizontal lines. Thus the explanation of discrimination is the same for both cases. In some sense this is saying that the question of how many lines are present is irrelevant, and that the important activity is instead to determine which terminators are present.

One reason for preferring this type of non-distinction between curves and lines is that there is little perceptual difference between a smooth arc and even a crude approximation to it by line segments. For example, considering regular polygons as texture elements, a hexagon is not distinguishable from a septagon. A similar finding is provided by Beck[1973] whose examples show that smooth arcs at a particular orientation have roughly the same discriminability as pieces of octagons at that same orientation (judged by speed of counting the number of differing element A within a texture of element B), and that small differences in the curvature of an arc produce only small differences in the discriminability.

Another possibly difficult subclass of textures to handle are those consisting of closed curves. One might think that overall orientation of some type might be needed here. But the combination of orientation and length of actual and virtual lines seems to behave




Figure 2.3.9

satisfactorily on many of these as well. For example, in fig 2.3.9 the actual line histograms differ. At 0 degrees, the outer has 1 at length 3 and 3 at length 1 while the inner has 2 at length 2 and 2 at length 1. At 90 degrees, the outer has 2 at length 2 and 2 at length 1 while the inner has 1 at length 3 and 3 at length 1. The same explanation would work if the corners were rounded. Note that as a closed curve becomes smoother, only differences on the basis of actual lines can be made so that if these agree, no discrimination due to shape will be predicted (as in the case of two similar blobs).

So virtual lines are connected between terminators within some local neighborhood of each other. The terminators are comprised of isolated points, endpoints of lines, and corners. This last category includes any sharp "join" of a line with a segment whether that segment is another line, an arc, or a line portion (e.g two lines crossing).

The above was a discussion of the necessity of the proposed theory. The next section will discuss its relation to the dipole theory of Julesz and the possibility of its sufficiency.

2.4 Relation to Julesz's theory

The evidence so far seems to indicate that the second order statistics limitation on texture discrimination (no discriminability if statistics agree) is sufficient to predict human performance. (There are no true counterexamples despite much searching by Julesz and the

author.) The operators proposed here are strictly less powerful; they might be said to compute "one and a half order" statistics. [Note 3] (First order statistics are not enough to handle virtual lines.) There are a number of ways in which the operators differ from dipole statistics.

The operators use a proper subset of the dipoles as input. For covered dipoles (corresponding to actual lines), all maximal dipoles are used. Maximal refers to any line delineated by terminators so that, for example, a plus contains 4 lines. For uncovered dipoles (corresponding to virtual lines), all local dipoles between terminators are considered. Note this eliminates non-maximal dipoles and allows only a finite number between adjacent lines.

The comparison is less fine. The operators consider lines grouped in equivalence classes of ± 10 degrees of orientation and compares a crude histogram of lengths within these; the dipole statistics require an exact match for both length and orientation. In addition, the operator difference must be sufficiently great to declare discrimination whereas any difference is crucial to the statistics (although Julesz seems to allow a bit of leeway in his informal discussion).

The locality restriction here differs somewhat from Julesz's. He considered only "separated" textures, those in which the elements do not overlap and are separated by some minimum distance r_0 . He implicitly made the assumption that $|r| < r_0$ for all dipoles r , thus considering dipoles only within elements, not those between them. (Gilbert and Shepp make this explicit.) Here virtual lines are drawn between terminators anywhere in a small neighborhood, irregardless of whether this contains more than one element. So virtual lines

will be drawn between elements if the elements are close enough together. This will often explain boundary effects (see section 2.6).

Some of Julesz's examples have different results under the experimental conditions proposed here. These include several sticky cases. For instance, fig 2.4.1, which Julesz[1973] considered a genuine counterexample, does *not* produce a discriminable region in 200 milliseconds (although it may appear slightly discriminable here).

Corresponding to the above differences, there are a variety of situations in which the proper prediction of no discrimination is made due to the the operators being correctly less sensitive than full second order statistics. These include the following 3 cases where two non-discriminable textures have the same actual and same virtual lines (with the meaning above) but differing dipole statistics:

(1) Any change in the lines which produces small length differences (e.g. less than doubling) and small orientation differences (e.g. less than 10 degrees). Such non-discriminable textures include a right angle versus a 95 degree angle, a hexagon versus a septagon, and "wiggles". This last situation can be illustrated by considering one texture element as a line and a point not on it and the second element as having the line with a small bump in it. The bump is not large enough to cause a terminator to be asserted but does cause the statistics to differ.

(2) Subthreshold differences such as in fig 1.3.9. Both the operators and the statistics differ but the operators must differ by more than a threshold to assert discrimination, which does not happen here. This is why Julesz's theory is not necessary; it is not specified how much the statistics must differ for discrimination to take place (and this specification may be very




Figure 2.4.1

difficult).

(3) The restriction of virtual lines to between terminators. In fig 2.4.2 the two textures are extremely indiscriminable. That the actual and virtual lines are the same can be seen by considering the construction of the generating elements; the basis is an X with a horizontal crossbar through the middle. The outer element has vertical lines connecting the upper left and lower right partitions while inner has vertical lines connecting the lower left and upper right. The elements are thus initially the same then two lines are added to each whose endpoints match endpoints of already existing lines. The actual line match is maintained and no new terminators are added; thus the virtual line match is maintained as well. The dipoles, however, do differ because dipoles exist between the vertical line on one side of an element and that on the other side; these dipoles slant to the left in the outer generating element and to the right in the inner. Those differing dipoles are excluded from the virtual lines considered since they do not occur between terminators (the terminators being taken up by the crossbars of the X and horizontal line).

The status of discrimination complexity may be summarized in the following table:

actual	virtual	dipole	discrimination	example	comment
+		+	+	1.1.1	first order (Marr)
-	+	+	+	2.2.5	
-	-	+	+	2.4.2	
		-	-	1.3.1	second order (Julesz)

Notes: if the actual lines differ, discrimination occurs regardless of what the virtual lines do. If the dipoles are the same, then the actual and virtual lines are necessarily the same as

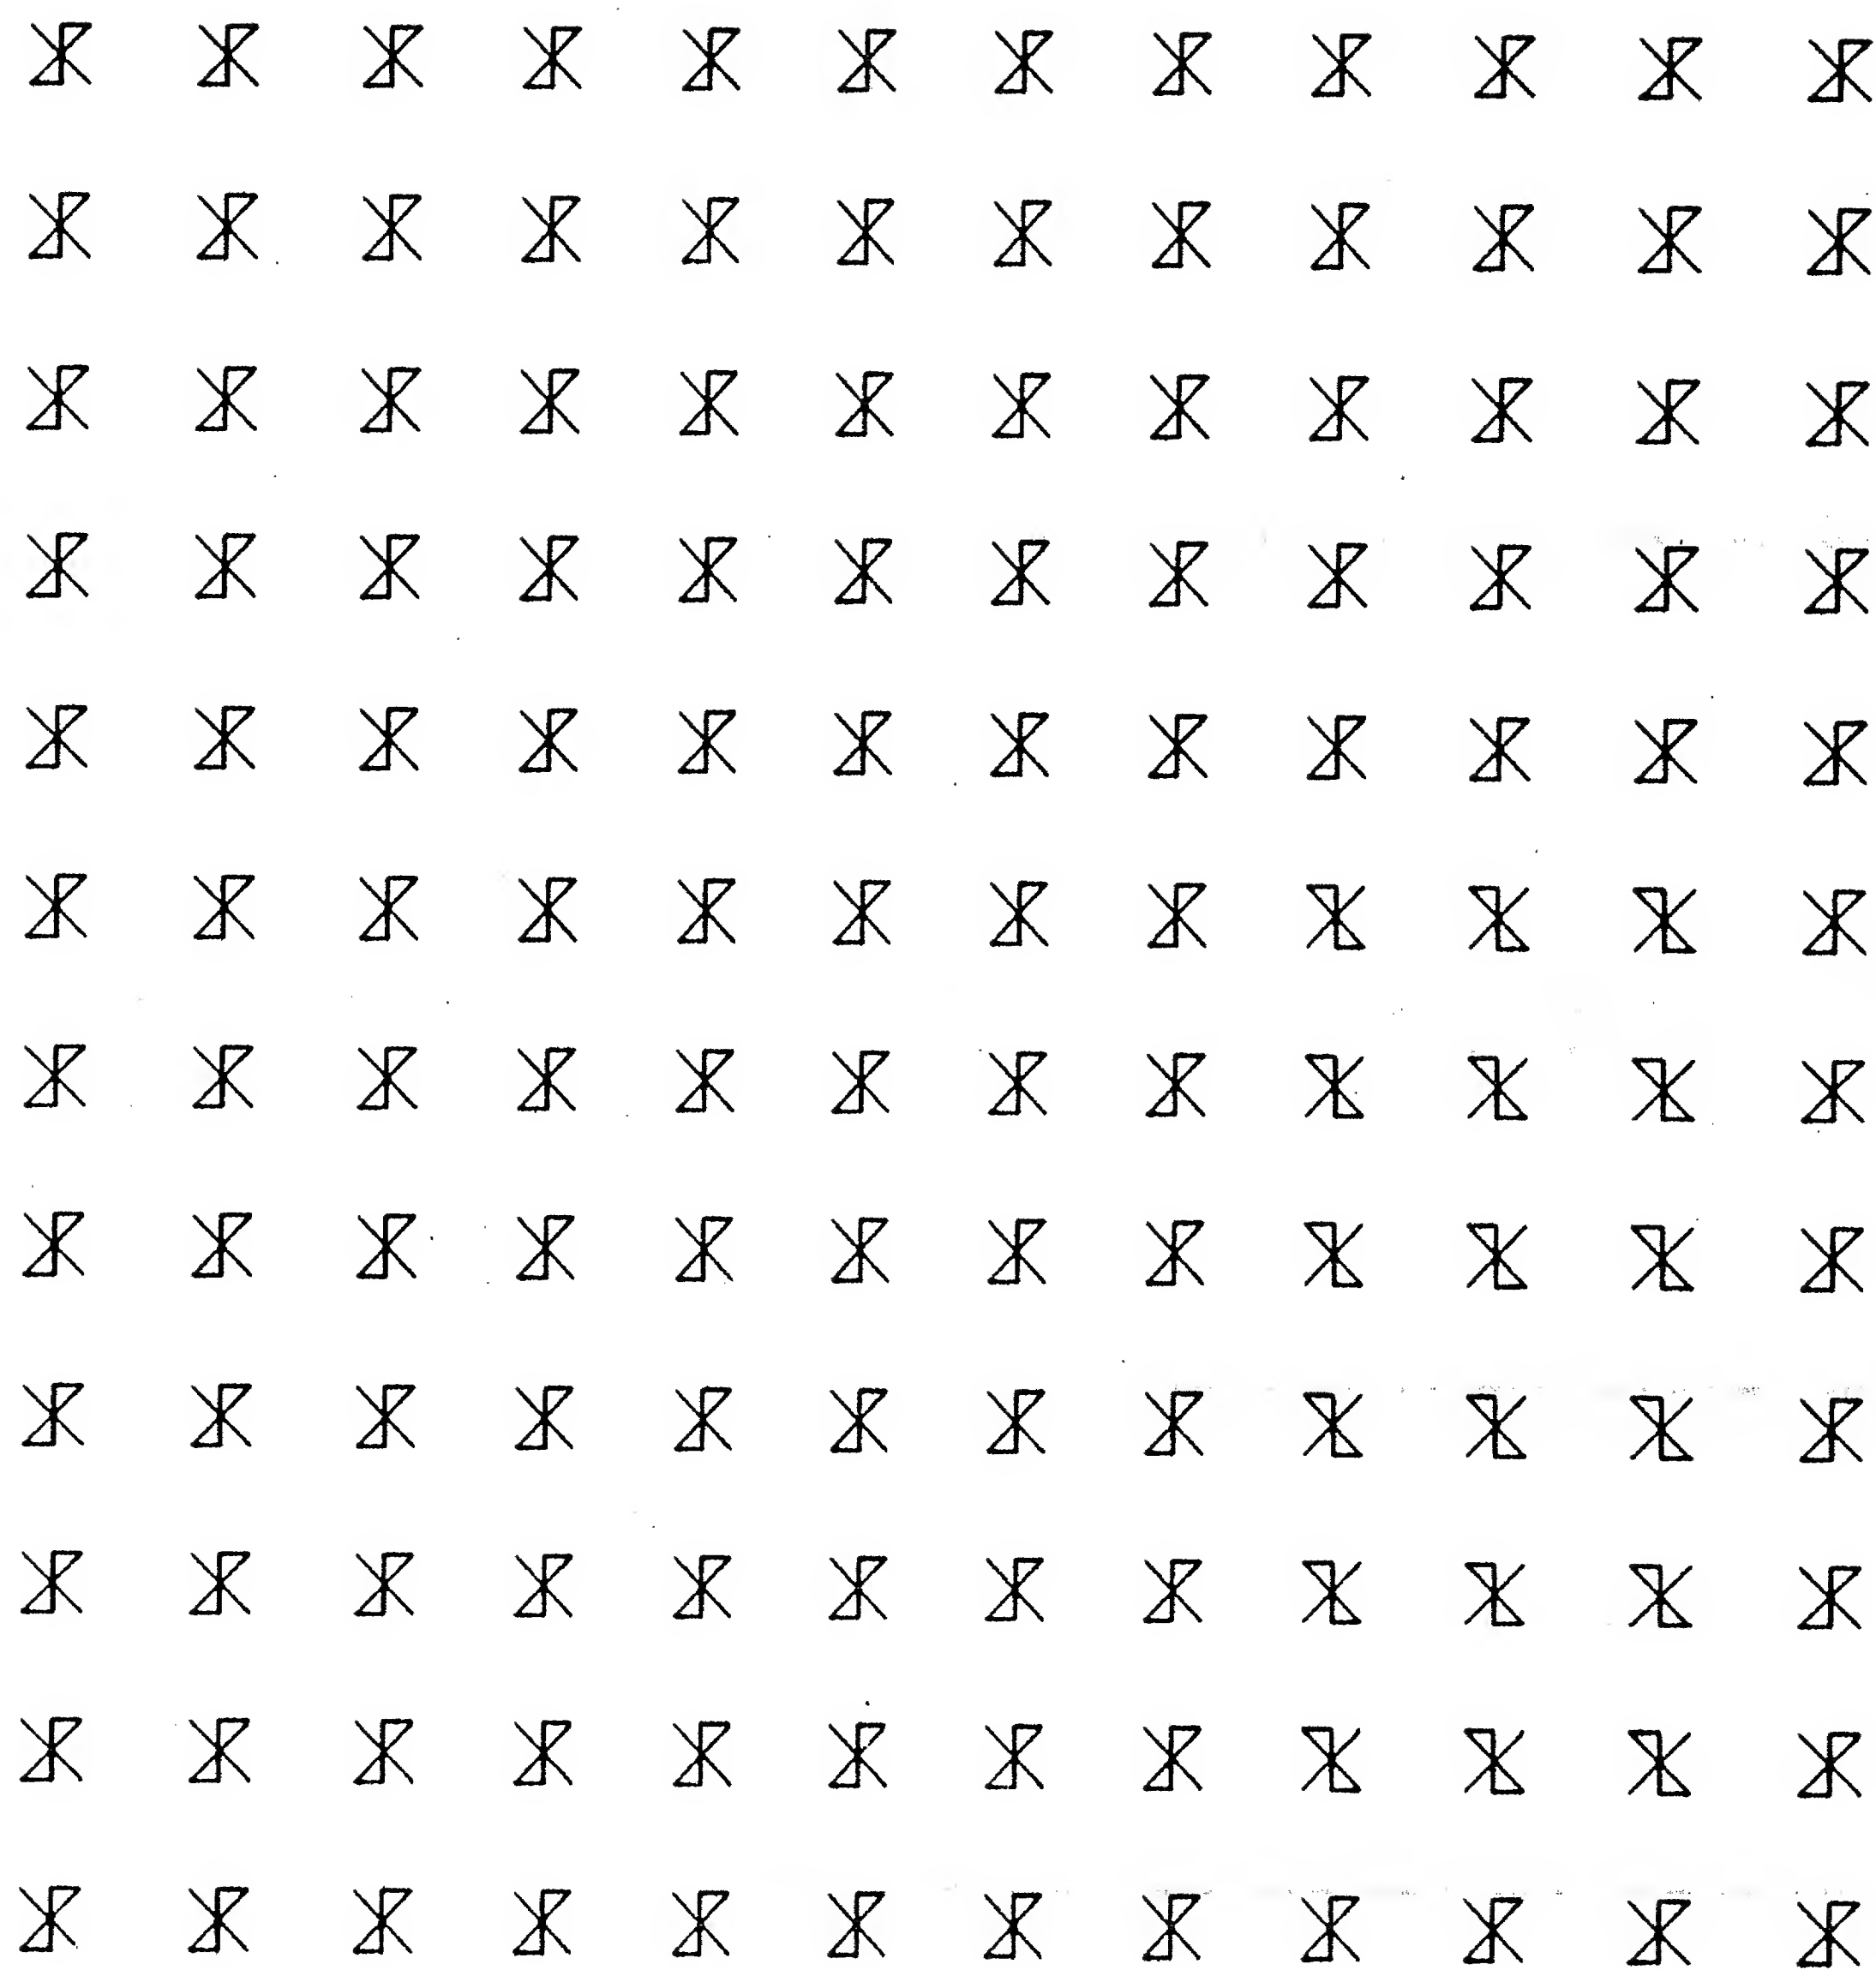


Figure 2.4.2

well.

As the above shows, the distinctions proposed here are less powerful than the second order statistics of Julesz. However, there is some reason to believe that they are still sufficient to predict human texture discrimination. It is conjectured that this is in fact true but it may be that further examples will require additional operators (c.f. sections 2.6 and 2.7). In any case, it appears that something less than full dipole statistics will suffice.

The next section surveys experimental results relevant to the proposed theory.

2.5 Psychological support


This section briefly considers data from the psychological and physiological literature which can be interpreted as supporting the proposed theory.

In addition to the previously mentioned conjecture of Julesz, there is considerable psychological evidence indicating that orientation, and in particular the type of orientation discussed above, is the primary factor in determining local shape. For example, Beck[1975,p406] concludes that "slope of lines is the most important of the grouping variables associated with shape". (He is considering similarity grouping, essentially element equivalence.) Julesz[1967] ran a multidimensional scaling test on textures formed from 2 by 2 arrays which indicated that brightness and orientation were the most important factors in discrimination. Olson and Attneave[1970] concluded that slope was a significantly stronger


factor than either arrangement or comparing straight versus curved segments. It also appears that the judged similarity of elements (Beck[1966]) or their familiarity (Julesz[1975]) are *not* strong criteria in discrimination.

Both operators do indeed seem to affect texture discrimination. It has been well-established that slope and arrangement are strong effects which can cause texture discrimination and which have different properties. Slope differences, elements with differing slope of lines, can be accounted for by differences in actual lines (operator (1)). Arrangement differences, elements with the same actual lines occurring in different positions, can be directly accounted for by the virtual lines (operator (2)). Thus the frequent references to arrangement in the literature could be replaced by references to differences in virtual lines (or at least predicted by these).

Virtual lines cause weaker discrimination than actual ones. This can be seen in fig 2.5.1 where the differences in both cases are caused by the diagonal lines which are virtual in (a) and actual in (b). There is also psychological data agreeing with this statement of the relative strengths of actual and virtual lines. The experiments of Beck[1967,1972], for instance, strongly indicate that elements differing only in arrangement are less discriminable than elements differing in slope (both in judged strength and length of time required to count one set of elements embedded in another). Fox and Mayhew[1977] reached a similar conclusion using reaction-time studies. Pickett[1970] pointed out that shape (actual lines) is more informative to humans than arrangement or density differences (virtual lines).



(b)



(a)

Figure 2.5.1

There is also neurophysiological and psychophysical data strongly suggesting the existence of orientation specific detectors. Orientation specificity of single cortical cells (selective firing to a limited range of orientation of bar stimuli) is well-established in a variety of animals including the monkey, whose visual system is quite similar to ours (Hubel and Wiesel[1968], Schiller, et.al.[1977]). It has been known for some time that humans appear to have orientation channels selectively sensitive to orientation with approximately ± 15 degree range. (Campbell and Kulikowski[1966]). Evidence for such channels has also been obtained using similar stimuli from human psychophysical data, human evoked potential responses, and cat single cell recording in the cortex. (Blakemore and Campbell[1969], Campbell and Maffei[1970], Campbell,et.al.[1969]). So humans appear selectively sensitive to orientation at a macroscopic level and probably this holds at a neural level as well. While orientation sensitivity could be needed for many activities, it is comforting to find that the properties essential for the early visual processing task of texture discrimination appear to be neurally implemented at low levels (e.g. in the cortex). [Note 4]

Thus far researchers recording from single cells appear to have only considered orientation sensitivity of physically present lines (bars). If the theory presented here is believed to accurately reflect low-level brain functioning, then one might be tempted to look for single cells at some early stage which are selectively sensitive to the orientation of *virtual* lines. Cells in area 17 probably cannot distinguish between actual and virtual lines due to small receptive field size but areas 18 and 19 are still a possibility. It would be interesting to see whether "lines" composed of three points or virtual lines between endpoints of an L have similar neural properties to actual lines.

2.6 Extensions

This section contains suggestions for extending the operators to less restricted domains; in particular, consequences of interaction effects and the allowance of other input primitives.

As texture consists basically of elements and gaps (Gibson[1950]), the major factors affecting it fall into two classes: (1) similarity and (2) proximity. Either of these can cause discrimination by itself. Similarity, the type considered in this thesis, might be termed "element equivalence". It deals with local shape comparisons between the texture elements with no consideration of the spacing between them. However, the spatial arrangement greatly affects discrimination. Proximity considers the relations (or their absence) between elements. It thus includes *boundary effects* where interactions between elements can cause discrimination between textures with similar elements (compare fig 2.6.1 to fig 2.6.2) and "spacing effects" where wide spacing forces lack of discrimination even with very dissimilar elements (consider the limiting case of elements several inches apart). As these occur even when there is constant density between two textures, the absolute spacing difference is the contributing factor.

When the elements become very close together, direct interactions between them, boundary effects, begin to play an important role in discrimination. Such phenomena as

T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T
T	T	T	T	T	T	T	T	T	T	T	T

Figure 2.6.1

333333333333
333333333333
333333333333
333333333333
333333333333
333333333333
333333333333
333333333333
333333333333
333333333333
333333333333
333333333333

Figure 2.6.3

subjective contours (fig 2.6.3) and accidental alignment (fig 2.6.1) can occur. These can often be detected in the computation here by the appearance of many virtual lines filled-in between elements which appear at the borders of the textures. Thus a large change in the length and orientation of virtual lines would predict discrimination in these cases. The present computation will handle these types of boundary effects as it fill in virtual lines within a neighborhood regardless of the number of elements within that neighborhood. As brought out previously, Julesz[1973] considered only "separated" textures with well-spaced elements and thus did not discuss proximity effects. There are also inverse boundary effects where narrow spacing causes discriminable elements to become non-discriminable textures.

The textures here consist of lines and points, essentially line drawings. However, real images have more complicated objects in them. As mentioned in section 1.2, Marr proposed a set of primitives which he conjectured were sufficient to represent all of the information contained in gray-level intensity images. Lines have been considered above; the following is a brief discussion concerning blobs and edges, the remaining primitives (Marr[1976,p497]).

A possible distinction among inputs is between sharp and fuzzy cases. The sharp case, where individual lines and points are distinguishable, was considered here. The fuzzy case consists of blobs, nebulous closed shapes. (Note this includes only smooth closed curves; the others can be approximated by lines.) It might be possible to handle blobs (and elements too fuzzy for individual lines to be distinguished) as follows. If the elements are large, discrimination can be accounted for by density changes while if small, change in the virtual

lines between them becomes a factor. This could explain the discrimination, for example, of elements formed by three circles in a horizontal row versus those formed by three in a vertical column.

Gradual sharpening is one way to think of the process of perceiving a progressively better image (Wohlfahrt in Woodworth[1938,p77]). First only formless shapes, blobs and bars, are seen, then as these become finer and sharper, points and lines appear. In the former coarse stage, density measures predominate in texture discrimination while in the latter, finer stage, local shape considerations begin to play a factor. Density changes can also catch some differing textures whose overall brightness is the same (as with small bright blobs versus large dim ones). [Note 5] The lengths and sizes of lines and blobs should also be weighted by the brightness.

Another distinction is between discrete and continuous intensity regions. The former have been considered in previous sections. The latter are possible when various shades of gray are permitted and the simple case of black areas sharply delineated by white may not occur. Here a principal factor is brightness clustering. (See, for example, Julesz's claim that discrimination is caused by clusters of proximate points of similar brightness [1971,section 4.4; also 1967]). This type of texture region, where the intensity varies smoothly, can be separated by finding edges between adjacent textures. Computational techniques for accomplishing this are well-known (e.g. Hueckel[1973], Marr[1976]).

2.7 Other possible operators

This section discusses several other reasonable operators and why they were not chosen. These include overall orientation and parallel lines.

It might have been noticed that the operators make no mention of texture elements. The claim has been that local shape is determined by orientation and length of actual and of local virtual lines. This suggests that no grouping into elements takes place at this stage of immediate texture region discrimination and thus that references to elements are an expository aid only and not a computational necessity. But other views are possible.

The orientation of actual lines must be used in any case. For example, fig 2.7.1 shows two discriminable textures generated by elements whose actual lines differ but whose virtual lines and overall orientation do not. But instead of using virtual lines between terminators to account for local shape in such cases as arrangement differences, one might use the overall orientation of the texture elements. For example, Marr[1976] proposed doing a preliminary grouping on the line and point assertions to produce a set of "place tokens" (presumably the texture elements here), then looking for boundaries by comparing the major axes of these. (Which could be found by a skeletonizing procedure, for instance. See Duda and Hart[1973, p356ff].) This might be contrasted as an "overall" orientation theory as opposed to the proposed "individual" orientation of (virtual) lines.

The overall theory is not preferable for several reasons. (a) Since it yields only the




Figure 2.7.1

orientation (and perhaps strength) of the major axis of an element, it does not provide a good mechanism for checking for a "sufficiently large" difference between elements. This could perhaps be crudely overcome by measuring the density of each element as well. (b) It is inherently an element equivalence theory and thus provides no insight into the important issue of handling boundary effects (relations between elements) which the virtual lines often handle automatically. (c) There are many cases, such as a plus, where the overall orientation is ambiguous. (d) As described in previous sections, there is substantial evidence that virtual lines are in fact used in texture and in other pieces of the vision processing computation.

Notice that both theories are two-part and a decision between them must be made by comparing textures with the same actual lines.

The evidence seems to favor the individual theory. There are a number of possible methods of finding the overall orientation and all seem to have at least mild counterexamples to proper prediction of discrimination. Consider the following methods:

- (1) Major axis of the convex hull. Even such discriminable elements as the cups in fig 1.3.2 have the same convex hull.
- (2) Major axis (unspecified) or Center of gravity line (the minimal length line segment y which minimizes $\sum d(x,y)$ for all points x in the element). It is possible to vary the spacing between the lines within a texture element in a way which does not change these overall orientation but does change the virtual lines. And in fig 2.7.2 this produces two elements which respectively generate textures that are discriminable.
- (3) Large masks (convolving the image with a mask roughly the size of an element and selecting the mask orientation with the greatest value (c.f. Marr[1976])). Consider




Figure 2.7.2

generating, on a uniform gray background, elements drawn by lines consisting of a black line/white line pair. A large mask should see only the average intensity of each element which is the same as the background, and thus the image should appear to be a uniform gray region. But in fact such textures are often discriminable.

If one then believes the individual line theory proposed here, it appears that immediate texture discrimination is a very low-level operation which occurs after retinal intensity processing but before even simple grouping. This claim is contrary to many other attempts at an explanation of texture processing and implies that the local shape of elements is only implicitly determined.

Another processor which is probably used in early visual processing is one that considers groups of parallel lines. Lines less than 10 degrees different in orientation are usually not distinguishable when compared texturally. However, groups of parallel lines are unusually strong and differentiable. Marr[1976,p507] and Stevens[1977] also make this point. Consider, for instance, fig 2.7.3 where discrimination occurs although the outer texture has lines at 90 degrees while the inner has intermixed ones at 82 and 98 degrees. The implementation here provides one explanation of this, namely that parallel lines all fall in one place in their orientation class (bucket) while even slightly non-parallel lines are distributed throughout theirs. Thus during comparison when the buckets are overlapped and their centers are shifted, ordinary lines can be separated and differentiated whereas parallel lines cannot. However, this important phenomenon may very well be necessary for discrimination in some cases and certainly merits further investigation.




Figure 2.7.3

Some sort of columnar "lining-up" could also be a possibility since perfectly regular textures appear in well defined columns so that slight deviations in size or shape or position stand out abnormally well. This was discounted in the examples considered here by testing textures with their centers slightly perturbed (randomly). That has the effect of decreasing the discriminability (since noise has been introduced into the windows being compared). However, it did not significantly change the class of textures considered discriminable (c.f. fig 2.7.4) although it did affect the parameter values.

Two shape descriptors not psychologically tenable are corners and crossed-lines. For instance, observing the placement of corners incorrectly predicts that the inverted cups (fig 1.3.1) is discriminable. Looking for crossed lines fails to predict that fig 2.2.3 (a), where the crossbar is three-quarters of the way up, is less discriminable than (b), where the crossbar is halfway up.

Other operators are possible, e.g symmetry, perpendicularity, connectedness. Pickett[1968] gives a survey of psychological evidence concerning a variety of factors. Duda and Hart [1973, chap 9] discuss a variety of more complex shape descriptors. However, the orientation and length of lines seems satisfactory so far for the crude type of local shape apparently required in immediate texture discrimination.

This chapter has discussed essential properties for determining texture discrimination and their embodiment as operators. Chapter 3 will be concerned with an implementation strategy for these that has a number of desirable computational and psychological properties.

Chapter 3

Computation of Texture Boundaries

3.1 Overview

The above analysis has provided a strategy for computing the boundaries of immediate texture regions. This chapter will be devoted to discussing details of implementing that strategy in a fashion which has desirable computational and psychological properties.

To compute texture discrimination, the first step is to account for the factors previously assumed as constant, i.e. overall brightness, color, motion, and the like. Then the density comparisons to catch many proximity changes are made. Finally one comes to the shape comparison which was the major topic of the previous chapter and will be the same here.

The basic procedure is to scan through the image looking for changes in the density of orientation. The scan is done by moving a window of "proper" size across the image and comparing the length and orientation of the lines in adjacent windows. Boundaries of texture regions occur when sufficiently large changes are found. Because the comparison difference must be sufficiently large, the difference between the windows is taken in a ratio against the total length of lines in the two windows. Thus texture discrimination occurs if and only if

$$\left[\frac{\sum \text{Difference}}{\sum \text{Total}} \right] > \text{threshold}$$

where difference means difference in length histograms summed over orientation and total means the sum total length of the lines in both windows. Two scans are made corresponding to the two texture operators: first check for differences based on actual lines. If no discrimination has occurred, fill in all local virtual lines between terminators and check for differences based on these.

The claim (c.f. Chapter 2) is that this procedure finds all texture boundaries in the line and point world, subject to the restrictions mentioned. It should be noted that no strong claim is made for the uniqueness of the computation. Other implementation techniques, while they seem unnecessarily indirect, should work successfully as long as they conform to the principles outlined previously.

Specific details of the computation will now be focused upon along with some notes on implementation of these.

3.2 Windows

What type of method should be chosen to search for variations in the orientation of lines? A window to scan over the image and compare adjacent values suggests itself as a convenient possibility. Image objects are affected by objects in a small neighborhood about

them in any case and some sort of averaging is desirable to prevent being fooled by noise and accidental effects. Windows are a natural suggestion when dealing with density and the notion of varying sizes of windows covering the visual field is reminiscent of the receptive fields of neurons in visual cortex. So how should these windows be manipulated?

Essentially what is wanted is to compare adjacent texture elements in adjacent windows. The window size is critical. If it is too big, an averaging effect takes place and no boundaries are found at all. If it is too small, adjacent windows contain substantial portions of the same elements and unwanted intraelement boundaries are found. Thus a fixed window size is unsatisfactory. The solution adopted here is the following. (No claim is made for the psychological reality of this. The visual system most likely combines its covering and overlapping array of neural receptive fields in fairly sophisticated ways.)

Start with a maximum window size. If the element size is larger than this, the system malfunctions (humans perform poorly on textures composed of large elements too). Scan the image with this window and mark all boundaries found. Shrink the window size slightly and repeat the scan. Iterate until there is a very large jump in the number of regions delimited by the boundaries. This presumably occurs because boundaries are now being found between an element and the gap next to it. Thus the approximate element and spacing size has been found. (Note this is much easier than actually grouping the elements into place tokens.) Let the boundary assertions which are made for this phase of region separation be those from the window size just before the blowup.

The windows in a particular scan also need to be overlapping. Otherwise, for example, half of a sharp change in orientation might fall into each of adjacent windows. So the

positions of the box centers should be varied slightly to avoid accidentally missing boundaries. The current implementation moves the windows separately to the right by one-half and up by one-half the window size to account for this. Also the windows are rectangular and scanned through only in the vertical orientation. (Other variations are possible.)

The spatial frequency component of textures could be calculated from the optimal window size. In the computation here, however, the initial density scan catches these. Note that while discrimination is insensitive to uniform density changes (doubling the size of both textures), density variations and spacing effects do make a difference. The spacing of the elements in particular deserves a bit of discussion. (It has been somewhat ignored in the argument so far which concentrated primarily on texture elements and their shape without regard for the important issue of spacing.)

There are essentially two possibilities for the character of spacing effects. The first is that spacing is a very crude effect on discrimination so that if elements are very close together, virtual lines occur between elements causing boundary effects, whereas if elements are very far apart they exceed the maximum window size and no discrimination occurs at all. In this view, any amount of spacing in-between is fairly similar so that one could vary the interelement distances quite a bit without affecting the discrimination as long as they did not get too close or too far apart. The second possibility is that the visual system is fairly sensitive to spacing and that discrimination is inversely proportional to density. In this view the closer the elements, the greater the discrimination. Note that the method

above of choosing the proper window size does not pay attention to the actual size chosen. Since the size is really roughly the element size plus half the spacing size, this implicitly assumes the first view. However, an experiment, varying the spacing of fig 2.6.1 from very close to very far, seemed to produce a continuum instead of a three state discrimination. So probably the window values need to be normalized for the window size.

Once the window size is decided, the next problem is to compare adjacent windows.

3.3 Comparison of windows

The assumption was made above that finding texture region boundaries is a local operation. The window size was selected with this in mind. Thus comparing adjacent windows and asserting a boundary point if they differ is probably a satisfactory strategy. The comparison, as mentioned previously, is on the length and orientation of lines in the windows. A basic point to remember is that immediate texture discrimination makes only very crude distinctions. Details of the comparison will now be examined.

The first observation is that orientations should be considered in equivalence classes. Lines fairly close in orientation to a given line should be considered "identical" to it. This can be accounted for by considering the lines as falling into buckets of ± 10 degrees width (the next section gives the rationale for choosing this number). Any line whose orientation

falls within a bucket is considered the same. The buckets should overlap to provide adequate comparison. Thus the decision was made to have buckets at 10 degree intervals with ± 10 degree ranges. (Note this means that every line falls into two buckets.) When comparing buckets from adjacent windows, the problem of possible missing of differences due to accidental placement again arises. (e.g. 9 should not differ from 14 degrees.) This can be overcome by varying the bucket centers over a range (e.g. every 10 degrees starting with 0 and then starting with 5) and recording a difference if any occurs. Also, as humans are particularly sensitive to horizontal and vertical lines, the 0 and 90 degree buckets are weighted slightly more heavily.

The difference between adjacent windows should be calculated by doing a subtraction, orientation bucket by bucket, of the length histograms in the adjacent windows. The lengths in these histograms are, as for the orientations, grouped into crude equivalence classes. For instance, normalize by the smallest number and consider only integral clumps (e.g. must be twice as big to be different). A line must be approximately 1.5 to 2 times longer than another to appear different from it in the texture case (Riley[1977]). Another possibility is to compare first the sum total of the lengths in a bucket, then the average length. In any case, only a crude histogram should be used.

As mentioned previously, the window comparison should be done by $(\Sigma \text{difference}) / (\Sigma \text{total}) > \text{threshold}$ as opposed to merely $\text{difference} > \text{threshold}$. One reason is that one would like to declare discrimination only in cases where there is a sufficient amount of difference and not just any (possibly small) amount. Another is that




Figure 3.3.3

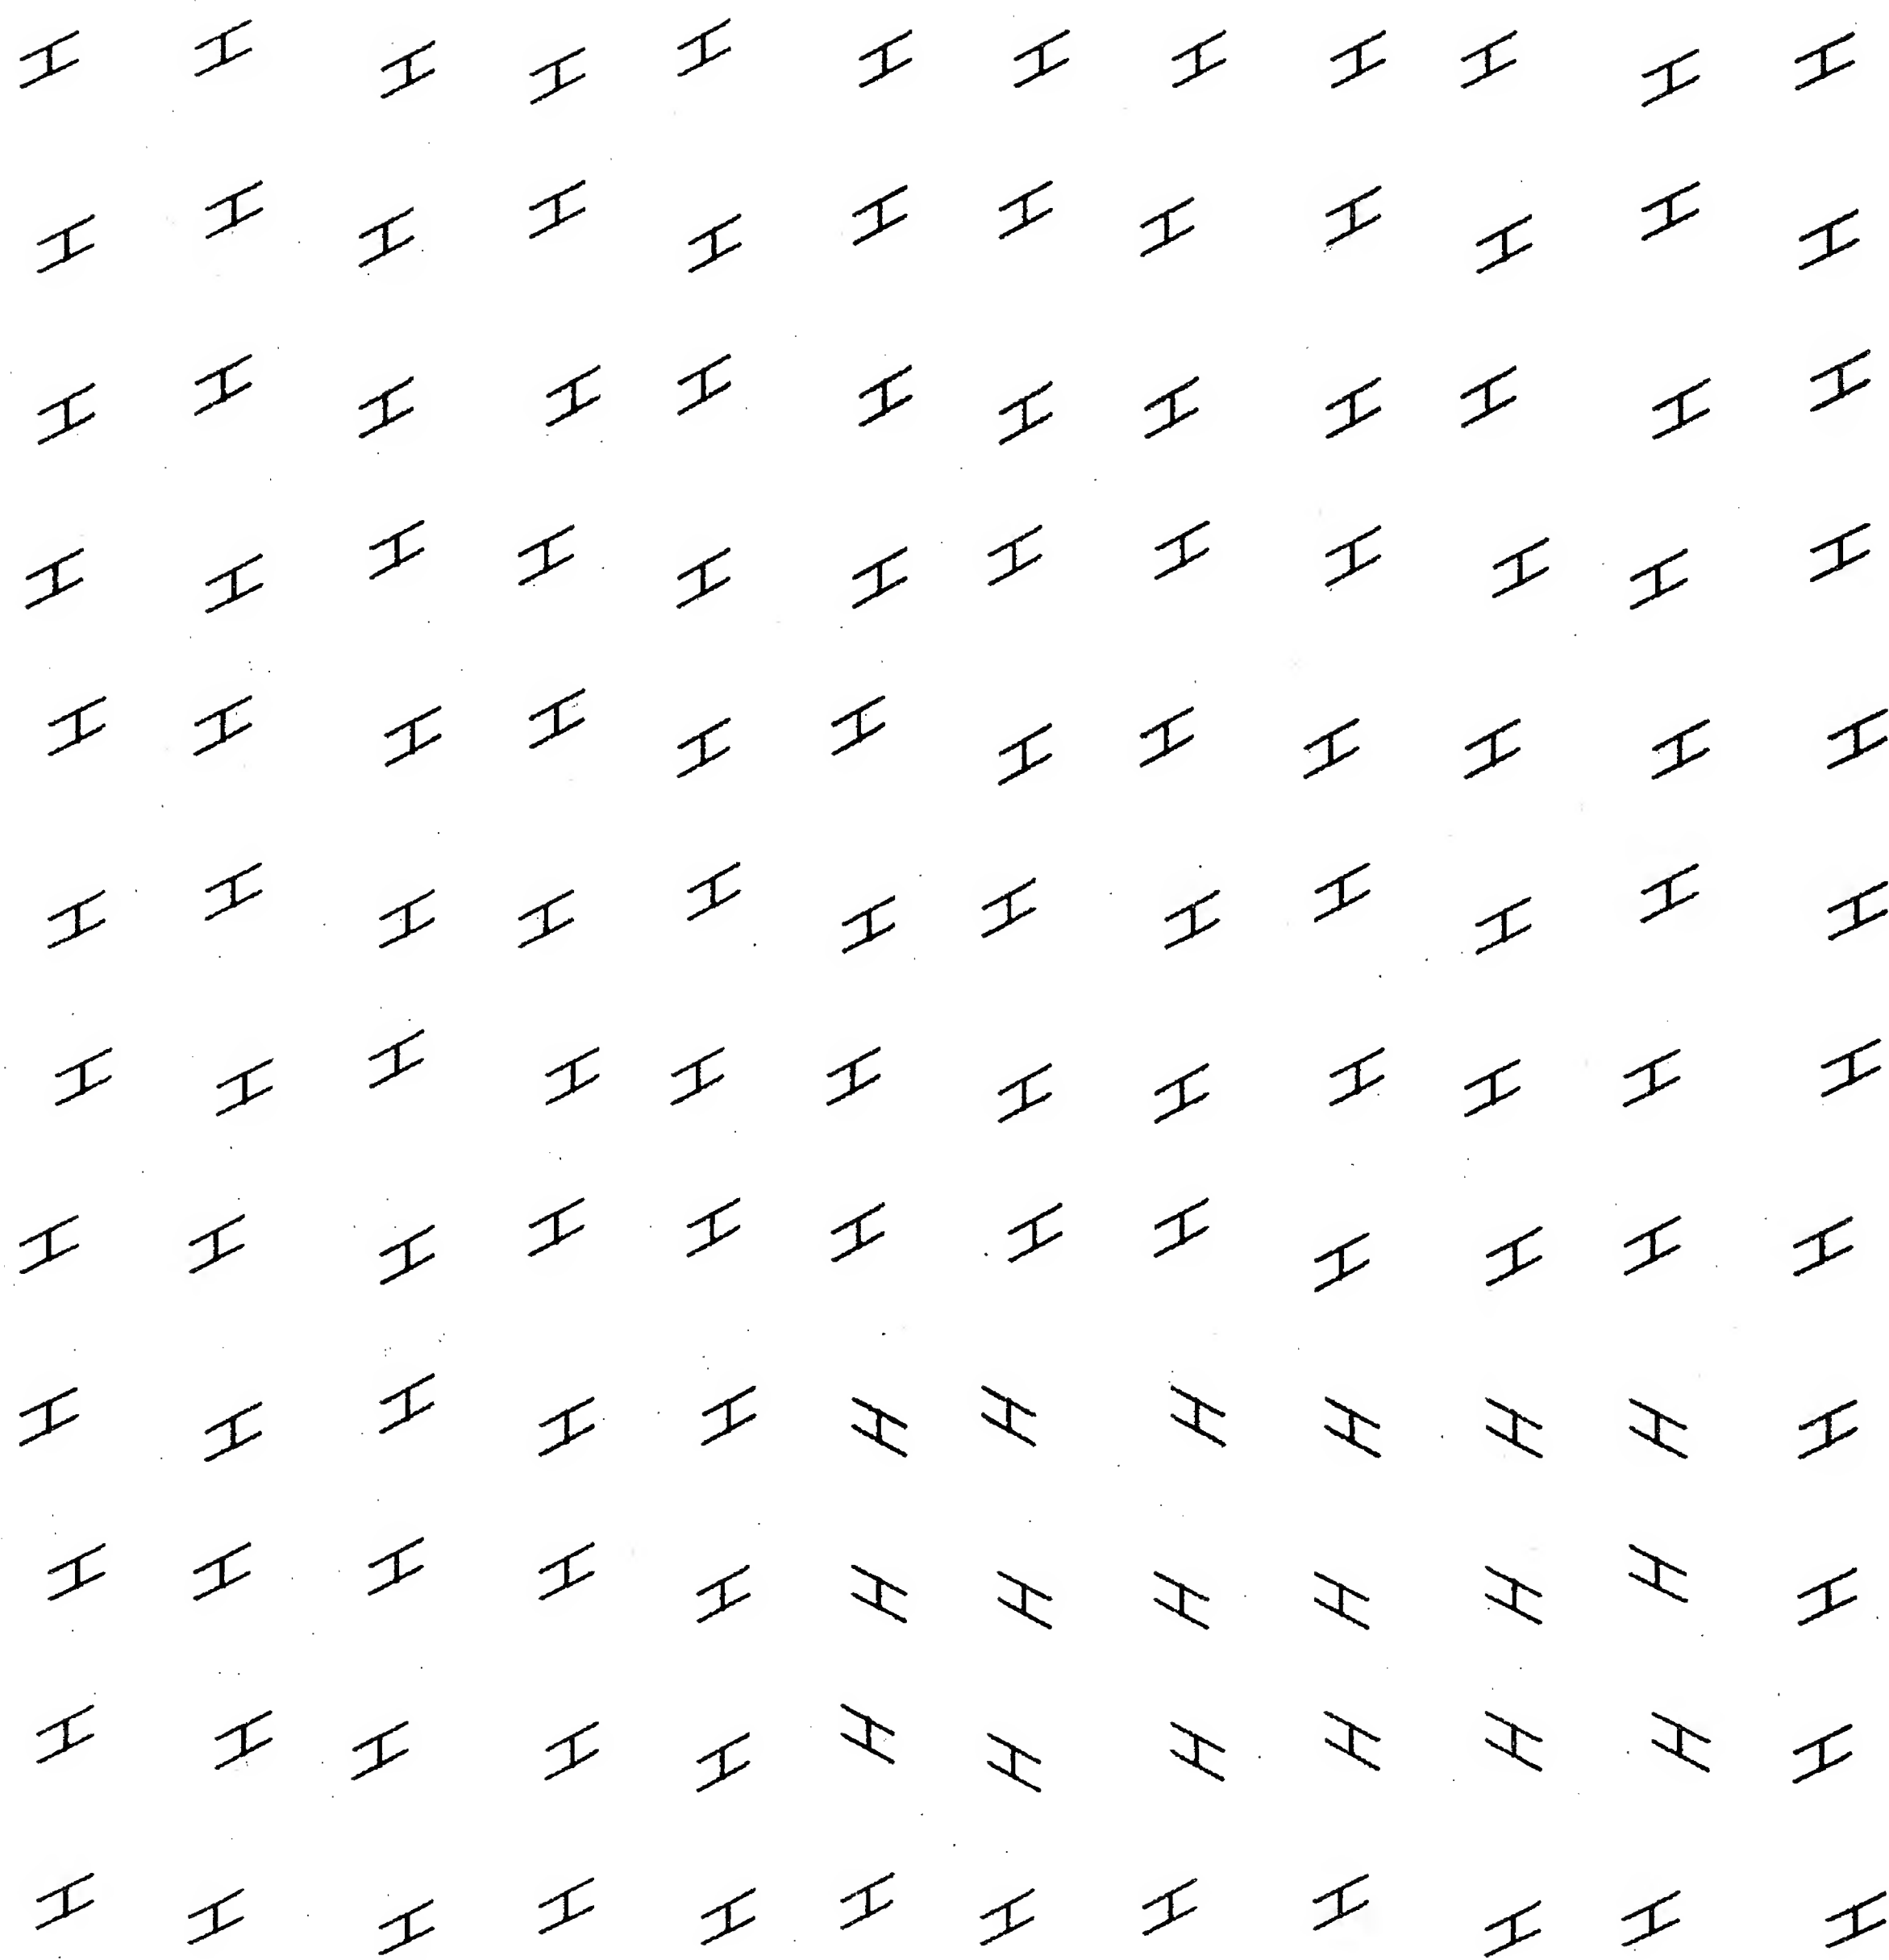


Figure 3.3.2




Figure 3.3.3

that the strength (or speed) of discrimination is proportional to the number of lines in the texture elements. One-line elements such as fig 1.1.1 produce much stronger differences than three-line elements such as fig 1.3.2 . This is true even if the orientation differences are not as great. In the limit, this effect produces the phenomenon of *clutter*, where even large differences in orientation are not perceived if the differences are small compared to the total lengths of lines. Thus in figs 3.3.1 to 3.3.3, discrimination goes down appreciably as the density of the elements increases. Virtual clutter also exists (c.f. fig 2.3.6) . A related effect is overload where, in the presence of too many lines, a skeletonized version of the image is processed. While this is likely an important property of the visual system, it has not yet been necessary to introduce it into the texture computation.


There is, however, an alternate possible method of comparing values from the orientation buckets. One could use a very gross measure, a two or three valued logic, and say that a bucket either has a line in it or not (i.e. 1 if any piece of any line whose orientation falls within the bucket range lies within the window, 0 otherwise. The 3-valued plan says none, some, or many.) The impetus for such a "gross" strategy comes from evidence that humans are remarkably poor at sensing variation in orientation. For example, in fig 3.3.4, 80% at one orientation and 20% at another is not differentiable from 50% at each. (Another example is that 3 orientations appear the same as random. Riley[1977] contains many such observations.)

How does this compare with the proposed difference/total or "difference ratio" method, i.e. taking the sum of the orientation differences in the two buckets over the total length of

INSIDE

BAR-LENGTH: 30.

ORIENTATION:



OUTSIDE

BAR-LENGTH: 30.

ORIENTATION:




Figure 3.3.4

lines in the two ? One reason to believe that differences must be considered in light of the total amount is to prevent the computation from predicting that a small notch on one of an indiscriminable pair of elements suddenly causes discrimination. Another support for the ratio method is that it provides a mechanism for considering length histograms as opposed to using a single number for each orientation bucket. As the gross method could not effectively manipulate a length histogram (which was shown to be desirable in section 2.2), it is less preferable. However, the ratio method has the disadvantage of requiring a threshold of some type if one wants to determine whether discrimination occurs or not. A threshold could be avoided in the gross method if any difference creates a difference. So it is worthwhile to consider whether psychological evidence favors one or the other.

One approach is to consider the handling of clutter (as above, lack of discrimination due to great density within elements). With the gross method, if the total is greater than some threshold, clutter and no discrimination occurs. Here "total" means the number of buckets with a 1 in them. If no clutter occurs, then any difference causes discrimination. (The threshold is probably 5 buckets as 3 orientations appear the same as random.) With the ratio method, the difference/total ratio is used directly and yields discrimination or not depending on the value. The no discrimination case, where the difference is small compared to the total lengths, includes clutter. Here total means the sum of the (weighted) line length histograms. (There could be a second threshold corresponding to good versus poor discrimination as well.) These two methods can be distinguished experimentally.

There are two deciding cases. Firstly, if the total and the difference are small, the gross theory can predict discrimination (no clutter since small total and some difference) while the

ratio theory can predict no discrimination (if the ratio of the small difference to the small total is small). An example of this appears in fig 3.3.5. The actual line ratio is .21 and the virtual line ratio .23, both below threshold. But only 3 buckets are used. So since no discrimination occurs, the ratio theory is favored. Secondly, if the total and difference are large, the gross theory can predict no discrimination (large total implies clutter) whereas the ratio theory can predict discrimination (the large difference forms a large proportion of the large total). An example of this appears in fig 3.3.6, where discrimination does indeed occur. There are 8 orientations (implying clutter) although the difference ratio for actual lines is .65. Thus the ratio theory seems to be preferred by humans. It should be noted that a difference ratio based on a small-valued logic would work as well and probably is fairly close to what is actually used.

The difference threshold value does not seem to be very critical. Using approximately 0.5 seemed to explain most cases. A careful determination of this value should be done. Note that there are actually two thresholds: one for actual lines and a slightly higher one for the slightly weaker virtual lines. (Perhaps the various types of virtual line should also be weighted differently by having each possess its own threshold.) Currently the actual and virtual line comparisons are done separately (the actual first due to its stronger nature). Another possibility would be to use a strategy where (a) if the actual lines differ, discrimination occurs and (b) if the actual do not differ but the virtual do, then discrimination occurs unless the actual are overwhelmingly similar. The latter rule can be considered as representing *competition*, where the actual and virtual lines are considered

I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I
I	I	I	I	I	I	I	I	I	I	I	I

Figure 3.3.5




Figure 3.3.6

simultaneously, weighting the former more heavily. If competition was incorporated into the computation, the comparison would be that discrimination occurs only when

$$\frac{(W * \text{actual_difference}) + \text{virtual_difference}}{(W * \text{actual_total}) + \text{virtual_total}} > \text{threshold}$$

where W, the weighting of the actual lines, is greater than 1.0 . No solid evidence for competition effects is known at present but it appears likely that this type of comparison is ultimately what should be used.

At some point the region boundaries found by the comparison process must be formed into regions. No special ideas on this have been investigated yet. Zucker[1976(b)] gives a survey of standard techniques. A brief digression into values of some of the parameters will be now made.

3.4 Parameter values

An attempt was made to empirically determine as many of the parameters as possible. This section discusses the computation of the virtual lines (primarily what the local neighbor distance is) and the width of the orientation buckets.

The local virtual lines are determined as follows. Pick a terminator, e.g. a line end or an isolated point. Find its nearest neighboring terminator (D away). Then draw in virtual lines to all terminators less than N times D away from the original point. N seems empirically to be 3.5 ± 1 .

N is in some sense the grouping limit number. Points within the $N \cdot D$ radius are the only ones considered for grouping to the original point although depending on the circumstances, some or none may actually be grouped. From several sources, it appears that this number is approximately 3. (Marr, Ullman, Stevens [personal communication] ; Atkinson, et al. [1976] who found that two sets of dots were grouped separately if separated by three times the interdot distance but not if separated by two times (measured by how many dots could be counted).) To attempt to confirm this, a direct test was (informally) run.

Glass and Perez [1973] observed that if a random dot pattern is rotated slightly and this superimposed on the original, a circular pattern appears. Similarly if the pattern is expanded slightly (x and y coordinates multiplied by a constant) and superimposed on the original, an exploded pattern is perceived. This led to the following experiment. Start with a random dot pattern. Since this original pattern is random, the nearest neighbor of any original dot in a superimposed pattern will usually be that same dot after the slight rotation or expansion. So if both a rotation and an expansion were superimposed on the original, three situations could occur:

(1) Circular pattern. (2) Random pattern. (3) Exploded pattern.

Presumably the perceptual mechanism involved is to check for each original dot O whether O -rotated or O -expanded is closer. This determines the nearest neighbor distance. If the

other distance is greater than N times this, only the nearest neighbor pattern would appear. If it is less, both patterns would appear, forming a random pattern composed of corners. So by varying the amount of rotation and expansion, all three types can be obtained and N can be measured by considering $R = \text{rotation distance} / \text{expansion distance}$. If $R < 1/N$, (1) should occur. If $R > N$, (3) should occur. Otherwise (2) should occur. This assumes that all dots were rotated or expanded a constant distance (as opposed to a constant degree of rotation or constant coordinate expansion).

Trying various distances ranging from 1 to 5 (with 2% density of dots) yielded the following. $R < 1/3$ produces a circular while $R > 4/1$ produces an exploded pattern. Intermediate R produce a random pattern. This leads one to suspect that N is between 3 and 4, which is consistent with previous results mentioned above. This grouping limit number may be a fundamental human information processing constant similar to the chunking number of Miller[1956].

The second parameter which seems fairly secure is the orientation bucket width. In examples such as fig 3.4.1 and others, discrimination seemed only to occur when line orientations were greater than 10 degrees apart (for both actual and virtual lines; they are 10 in the figure here). Campbell[e.g. 1966] finds the orientation by which one can change a sinusoidal grating without varying one's adaptation to it to be 12-15 degrees. Riley[1977] in a detailed quantitative study of human perception of orientation of lines concluded that the orientation equivalence classes were somewhere in the range of 10-30 degrees. So ± 10 degrees seems to be a convenient reasonable range. Another confirmation of this is that the smallest




Figure 3.4.1




Figure 3.4.2

n for which an n-sided regular polygon appears identical to an n+1-sided (in the texture case) is 6 and this is the first case in which the angles are less than 10 degrees apart (hexagon: 120 septagon: 128.6). See fig 3.4.2 .

A preliminary computer program embodying the implementation principles outlined above has so far produced reasonable results on simple straight line textures (although it has not yet been tried on any non-obvious cases).

3.5 Conclusion and future prospects

This thesis has discussed the early visual processing problem of immediate region discrimination based on texture information, primarily in the simple case of elements composed of lines and points. A case has been made for the usefulness of considering changes in length and orientation of actual and of local virtual lines to find texture boundaries. As this has been mainly an preliminary exploration, there are a number of aspects which need more work.

On the psychological end, a full, rigorous experiment needs to be run, using the experimental procedure outlined in section 1.3, to benchmark precisely what discriminations can and cannot be made in 200 milliseconds. As there is strong evidence that texture is predominantly a peripheral phenomenon (e.g. Beck[1972]), the textures should be examined peripherally as well as foveally as was done here. (The angle of viewing also seems to

affect discrimination.) Element equivalence tests, where large but isolated pieces of separate textures are compared, would help eliminate the possibility of boundary effects when attempting to discover what affects the shape of texture elements. The precise interplay between orientation and length and between actual and virtual lines needs to be determined. Parameter values such as the difference thresholds should also be empirically set. It would be nice to find definite examples settling the existence of virtual lines in general and virtual lines between corners in particular. More textures need to be examined to determine if other operators such as parallel lines, density, or even overall orientation are necessary.

On the formal end, the definition of line and terminator needs to be specified more exactly, especially in regard to corners and smooth curves.

The implementation strategy still has many unsettled points, as referred to throughout this chapter. For example, what exactly should go into the histograms and how should they be compared? Experimentation, both psychological and computational, needs to be done to decide proper ways of doing these. Extensions to less restricted domains, as mentioned in section 2.6, should also be considered. Finally, work should be continued to develop the program to a stage where it could be tried on real world textures (samples processed by the primal sketch). It will be interesting to see what additional factors are needed to explain the immediate discrimination of texture regions.

Notes

- (1) This thesis contains no formal mathematics. For pedagogical reasons, terms such as necessary, sufficient, subset, and so on will be occasionally used. However, they are not meant in the technical sense and should be considered only as heuristic guides. For example, "necessary" is used here as a psychological term where "property A is necessary" means that two textures can be exhibited whose discrimination seems to be solely due to property A. Similarly, "this set of properties is sufficient" will mean that a computation embodying them will predict discrimination or non-discrimination when immediately viewing textures in exactly the same cases that humans do. In a like manner, "upper bound" and "lower bound" refer to how much or how little is apparently psychologically necessary to implement the texture computation. This lack of formal proof means, of course, that further examples may show the assertions here to be incorrect.
- (2) The words "line" and "line segment" will refer to maximal physically present line segments. Given a line segment, there are an infinite number of subsegments contained within it, e.g. one starting at one endpoint and extending halfway to the other. These inner segments will be referred to as subsegments. The length of the smallest maximal segment will nearly always be assumed to be 1. Similarly, "point" will mean black point.
- (3) While discussion of order statistics is helpful to show the relation of this theory to previous ones, the author does not feel that statistical considerations are the best way of approaching the foundations of texture discrimination. An explanation of the processing involved seems more useful than a phenomenological description. This requires investigating the essential properties which determine discrimination in terms of, say, orientation of lines. Statistical considerations are very suggestive but do not seem to provide a detailed enough explanation. The author believes that the computational theory outlined here can be extended into a detailed enough form to test it with further psychophysical predictions.
- (4) It is certainly incorrect to say that the texture computation occurs at the cortical level. However, if one believes that the processing described here in fact has some relation to the processing in the brain, such an elementary and fundamental computation as texture discrimination must occur at a very low level.
- (5) A density processor is needed in any case as density changes can cause strong discrimination effects. Both density of points and of lines are crucial. For the former, humans appear sensitive to 20% changes in density of dots (c.f. Pickett[1970]). For the latter, consider the discriminable textures in fig 2.2.4 where the dot density is the same in both textures but the line density (number of lines per unit area) is twice as great in the inner.

References

- Atkinson, J., Francis, M.R., and Campbell, F.W., 1976. The dependence of the visual numerosity limit on orientation, colour, and grouping in the stimulus. *Perception* 5: 335-342.
- Attneave, F., 1974. Apparent motion and the what-where connection. *Psychologia* 17:108-120.
- Bajcsy, R., 1973. Computer description of textured surfaces. *Proc. 3rd International Joint Conf. on Artificial Intelligence*, 572-577.
- Beck, J., 1966. Effect of orientation and of shape similarity on perceptual grouping. *Perception and Psychophysics* 1:300-302.
- Beck, J., 1967. Perceptual grouping produced by line figures. *Perception and Psychophysics* 2: 491-495.
- Beck, J., 1972. Similarity grouping and peripheral discriminability under uncertainty. *American J. of Psychology* 85: 1-19.
- Beck, J., 1973. Similarity grouping of curves. *Perceptual & Motor Skills* 36: 1331-1341.
- Beck, J., 1975. The relation between similarity grouping and perceptual constancy. *American J. of Psychology* 88: 397-409.
- Blakemore, C. and Campbell, F.W., 1969. On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *J. Physiology* 203: 237-260.
- Brodatz, P., 1966. Textures: A Photographic Album for Artists and Designers. (New York: Dover Publications).
- Campbell, F.W. and Kulikowski, J.J., 1966. Orientational selectivity of the human visual system. *J. Physiology* 187: 437-445.
- Campbell, F.W., Cooper, G.F., and Enroth-Cugell, C., 1969. The spatial selectivity of the visual cells of the cat. *J. Physiology* 203: 223-235.
- Campbell, F.W. and Maffei, L., 1970. Electrophysiological evidence for the existence of orientation and size detectors in the human visual system. *J. Physiology* 207: 635-652.
- Duda, R.O. and Hart, P.E., 1973. Pattern Classification and Scene Analysis. (New York: John Wiley & Sons).

- Fox, J. and Mayhew, J.E.W., 1977. Region suppression sheds light on information - processing in texture discrimination. to appear.
- Gilbert, E.N. and Shepp, L.A., 1974. Textures for discrimination experiments. Bell Laboratories memo TM-74-1218-6 (April).
- Gibson, J.J., 1950. The Perception of the Visual World. (Boston: Houghton Mifflin).
- Glass, L. and Perez, R., 1973. Perception of random dot interference patterns. *Nature* 248: 360-362.
- Haralick, R.M., Shanmugan, K., and Dinstein, I., 1973. Textural features for image classification. *IEEE Trans. Systems, Man, & Cybernetics* SMC-3: 610-621.
- Hawkins, J.K., 1970. Textural properties for pattern recognition. pp 347-369 in Lipkin and Rosenfeld [1970].
- Hubel, D.H. and Wiesel, T.N., 1968. Receptive fields and functional architecture of monkey striate cortex. *J. Physiology* 195: 215-243.
- Hueckel, M.H., 1973. An operator which recognizes edges and lines. *J. Assoc. Computing Machinery* 20: 634-647.
- Julesz, B., 1962. Visual pattern discrimination. *IRE Trans. Information Theory* IT-8: 84-92.
- Julesz, B., 1965. Texture and visual perception. *Scientific American* 212: 38-48 (Feb).
- Julesz, B., 1967. Some recent studies in vision relevant to form perception. pp136-154 in W. Wathen-Dunn (ed), Models for the Perception of Speech and Visual Form (Cambridge: MIT Press).
- Julesz, B., 1969. Cluster formation at various perceptual levels. pp297-315 in S. Watanabe (ed), Methodologies of Pattern Recognition, (New York: Academic Press).
- Julesz, B., 1971. Foundations of Cyclopean Perception. (U. Chicago Press).
- Julesz, B., 1975. Experiments in the visual perception of texture. *Scientific American* 232:34-43 (April).
- Julesz, B., Gilbert, E.N., Shepp, L.A., and Frisch, H.L., 1973. Inability of humans to discriminate between visual textures that agree in second-order statistics -- revisited. *Perception* 2:391-405.

Koffka, K., 1935. Principles of Gestalt Psychology. (New York: Harcourt, Brace, & World).

Lieberman, L.I., 1974. Computer recognition and description of natural scenes. unpublished PhD thesis, U. Pennsylvania Dept of Computer & Information Sciences.

Lipkin, B.S., 1969. Psychopictorics and pattern recognition. pp 25-40 in Pattern Recognition Studies, Soc. Photo-Optical Instrumentation Engineers Proc., Vol. 18.

Lipkin, B.S. and Rosenfeld, A (eds), 1970. Picture Processing and Psychopictorics (New York: Academic Press). section III on Texture.

Marr, D., 1976. Early processing of visual information. *Philosophical Trans. Royal Society London (Series B)* 275: 483-524. also MIT AI Lab Memo 340.

Miller, G.A., 1956. The magical number seven plus or minus two: some limits on our capacities for processing information. *Psychological Review* 63: 81-92.

Olson, R.K. and Attneave, F., 1970. What variables produce similarity grouping? *American J. of Psychology* 83: 1-21.

Pickett, R.M., 1968. Perceiving visual texture: a literature survey. Aerospace Medical Research Lab AMRL-TR-68-12, Wright-Patterson Air Force Base, Ohio.

Pickett, R.M., 1970. Visual analyses of texture in the detection and recognition of objects. pp 289-308 in Lipkin and Rosenfeld [1970].

Ratliff, F., 1972. Contour and contrast. *Scientific American* 226: 90-101 (June).

Riley, M., 1977. Discrimination of bar textures with differing orientation and length distributions. unpublished BS thesis, MIT Electrical Engineering & Computer Science Dept.

Riseman, E.M. and Arbib, M.A., 1977. Survey: Computational techniques in the visual segmentation of images. *Computer Graphics & Image Processing* 6: 221-276.

Rosenfeld, A. and Thurston, M., 1971. Edge and curve detection for visual scene analysis. *IEEE Trans. on Computers* C20:562-569.

Schiller, P.H., Finlay, B.L. and Volman, S.F., 1977. Quantitative studies of single cell properties in monkey striate cortex. esp II. Orientation specificity and ocular dominance. *J. Neurophysiology* (to appear).

Stevens, K., 1977. Computation of locally-parallel groupings. MIT Artificial Intelligence Lab Memo 392.

Ullman, S., 1977. The Interpretation of Visual Motion. MIT PhD thesis. to appear as a Technical Report, MIT Artificial Intelligence Lab.

Woodworth, R.S., 1938. Experimental Psychology. (New York: Holt).

Zucker, S., 1976(a). On the structure of texture. *Perception* 5: 419-436.

Zucker, S., 1976(b). Region growing: childhood and adolescence. *Computer Graphics and Image Processing* 5: 382-399.

